Rearrangement Planning via Heuristic Search

Jennifer E. King, Siddhartha S. Srinivasa
Carnegie Mellon University (CMU)
{jeking, ss5}@andrew.cmu.edu

arXiv:1603.08642v1 [cs.RO] 29 Mar 2016

***Abstract*—We present a method to apply heuristic search algorithms to solve *rearrangement planning by pushing* problems. In these problems, a robot must push an object through clutter to achieve a goal. To do this, we exploit the fact that contact with objects in the environment is critical to goal achievement. We dynamically generate goal-directed primitives that *create* and *maintain* contact between robot and object at each state expansion during the search. These primitives focus exploration toward critical areas of state-space, providing tractability to the high-dimensional planning problem. We demonstrate that the use of these primitives, combined with an informative yet simple to compute heuristic, improves success rate when compared to a planner that uses only primitives formed from discretizing the robot's action space. In addition, we show our planner outperforms RRT-based approaches by producing shorter paths faster. We demonstrate our algorithm both in simulation and on a 7-DOF arm pushing objects on a table.**

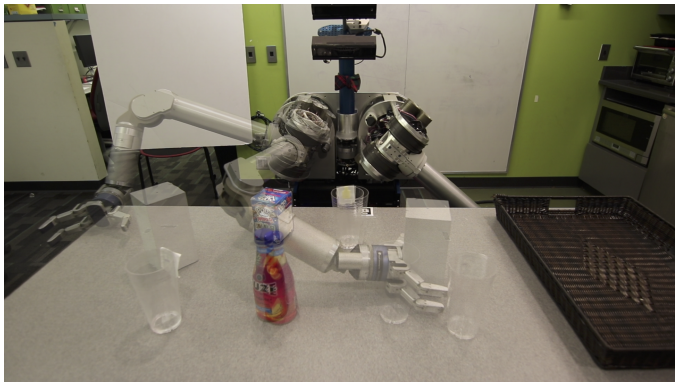

Fig. 1: A solution to a rearrangement problem. The robot moves the box across the table to make it reachable by the left arm. The robot freely interacts with and moves clutter in the scene in order to achieve the goal.

I. Introduction

In this paper we present a method for planning pushing actions that allow a robot to move an object to a goal pose through clutter. Consider the scene in Fig.1. Here the robot is tasked with moving the white block from the right side of the table to the left side of the table in order to make it accessible to the left arm to lift and place the block on the tray. All objects on the table can be moved by the robot and the final pose of these objects does not matter.

Rearrangement planning [4, 9, 25, 32–34] problems such as this are difficult for two reasons. First, the planner must search the high-dimensional state space containing the state of the robot and all the objects cluttering the scene.

Second, the system is underactuated and non-holonomically constrained. In particular, the objects in the environment can only move when contacted by the robot, and the motion of a contacted object is directly governed by the physics of the contact.

As a consequence, solving the two-point boundary value problem (2PTBVP) — connecting two states via feasible control inputs — is often analytically intractable and numerically expensive. This is in contrast with holonomically constrained systems, where often straight lines (more technically, geodesics in configuration space) trivially solve the 2PTBVP.

Randomized planning methods such as the Kinodynamic Rapidly-Exploring Random Tree (RRTs) [17] have been shown to be quite effective at handling high-dimensional state spaces and have proven applicable to the rearrangement problem [14, 16, 36] by never needing to solve the 2PTBVP exactly, instead sampling and rolling out actions and growing trees that never have to meet at a point. However, the resulting solutions are often quite suboptimal. Post-processing methods such as shortcutting [13, 29, 30] can be applied to improve the local optimality of solutions, but these solutions may still be far from the global optimal.

Proposed variants to the RRT algorithm allow the planner to find globally optimal solutions over time [10, 11, 15, 35]. But these methods depend on the ability to solve the 2PTBVP in order to "rewire" suboptimal paths through the graph.

Likewise, graph-based methods such as A* search [12] are attractive because they find globally optimal solutions when applied to discrete state and action spaces. Prior works have shown their applicability to continuous state and action spaces by first discretizing the state space, then connecting the discrete states with feasible actions [6, 7, 18]. However, even these methods require a solution to the 2PTBVP for creating the lattice.

Our first key insight is that, although we cannot solve the 2PTBVP in the full state space, we can solve it in the lower dimensional subspace containing only the robot. In other words, we can generate actions that move the robot

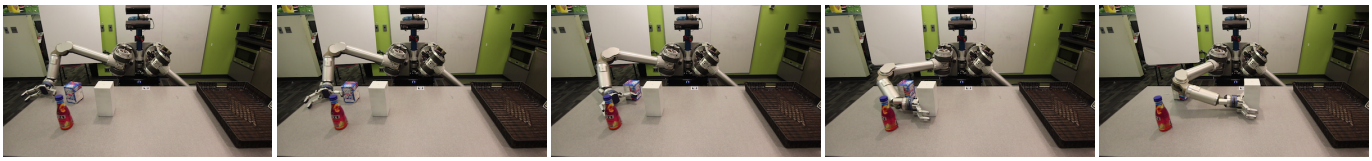

Fig. 2: A rearrangement planning problem. The robot is tasked with pushing the white box near the tray. The plan exhibits *contact* and *pushing* primitives to move clutter and achieve the goal.

between two configurations. This allows us to generate primitives by specifying desired configurations for the robot.

However, we are faced with two more challenges:

1) **Action space:** Typical lattice methods define a single set of actions, or primitives, to be applied at every state. Contact between robot and objects is critical to success in rearrangement planning and the set of primitives that create contact between robot and objects varies with state. This makes it difficult to define a single set of primitives that is useful across all states.
2) **Graph resolution:** Discretizing the continuous state space allows the planning algorithm to consider nearby states equivalent. This relies on the assumption that two nearby states will remain nearby when a primitive is applied. The discontinuous nature of pushing interaction means this assumption often does not hold. Even a very small change in the initial pose of an object can lead to very different final poses of the object after being pushed. Thus, applying naive discretization methods can lead us to incorrectly define two states equivalent, causing the search to miss important areas of state space.

Our second key insight is that we can dynamically generate primitives that create contact with objects at each state expansion during the search. We can use these dynamically generated primitives to supplement a core primitive set applied to every state. This focuses exploration toward critical areas of state space. The focused search, in turn, allows us to use a very fine graph resolution.

To dynamically generate primitives we exploit two further aspects of the problem: (1) we must *create* contact and (2) in the quasistatic environments we consider, we must *maintain* contact in order to move an object. We define *contact* primitives as those that move the robot to a configuration in or near contact with an object. Then, we use simple physics assumptions to generate *pushing* primitives that move the robot in a direction likely to maintain contact with an object.

The remainder of this paper is organized as follows. In Section II we define the rearrangement planning problem. We then define a core set of *basic* primitives, and the dynamically generated *contact* and *pushing* primitives in Section III. In addition, we show how these primitives, combined with an informative yet simple to compute heuristic, can be used in discrete search. We present experimental results in Section IV that demonstrate that the use of *contact* and *pushing* primitives improves success rate of a search based planner when compared to one that uses only *basic* primitives. In addition, we show that our planner produces shorter paths faster when compared to RRT-based approaches for solving the same problem.

II. The Rearrangement Planning Problem

Assume we have a robot, $\mathcal{R}$, working in a bounded world populated with a set of objects $\mathcal{M}$ that the robot is allowed to manipulate and a set of obstacles $\mathcal{O}$ which the robot is forbidden to contact. The robot is endowed with configuration space $\mathcal{C}^R$ and each object in $\mathcal{M}$ is endowed with configuration space $\mathcal{C}^i$ for $i = 1 \dots n$.

Our search is defined on the state-space $\mathcal{X}$ that is the Cartesian product space of the configuration spaces of the robot and the objects in $\mathcal{M}$: $\mathcal{X} = \mathcal{C}^R \times \mathcal{C}^1 \times \cdots \times \mathcal{C}^m$. A state $x \in \mathcal{X}$ is defined by $x = (q, o^1, \dots, o^m), q \in \mathcal{C}^R, o^i \in \mathcal{C}^i \ \forall i$. We define the free state space $\mathcal{X}_{free} \subseteq \mathcal{X}$ as the set of all states where the robot and objects are not penetrating themselves, the obstacles or each other. We allow contact between robot and objects in all states in $\mathcal{X}_{free}$. This contact is critical for manipulation.

The state x evolves nonlinearly based on the physics of the manipulation. The motion of the objects is governed by the contact between the objects and the manipulator. We describe this evolution as a non-holonomic constraint:

$$\dot{x} = f(x, u) \tag{1}$$

where $u \in \mathcal{U}$ is an instantaneous control input. The function f encodes the physics of the environment.

We define a primitive, a, as a discrete set of control inputs:

$$a = \{(u, d)_i | u \in \mathcal{U}, d \in \mathbb{R}^+, i = 1 \dots p\} \tag{2}$$

where u defines a control and d defines the duration to apply the control.

We identify a single object, $g \in \mathcal{M}$ as the *goal object*. We identify a goal region $\mathcal{X}_G$ as the set of states where

where the goal object is within radius r_{goal} of a desired configuration $p_{goal} \in C^g$:

$$\mathcal{X}_G = \{x | x \in \mathcal{X}_{free}, o^g \in x, \|o^g - p_{goal}\| \leq r_{goal}\} \quad (3)$$

The task of rearrangement planning is to find a sequence of primitives (path), $\pi := \{a_1 \dots a_t\}$, such that when applied to a start state $x_0 \in \mathcal{X}_{free}$ under the constraint f we end in the goal region.

We define the cost of a path, π, as the distance the end-effector of the robot moves in the configuration space of the goal object. Formally, assume we have a distance metric, $\mathrm{d} : \mathcal{C}^g \times \mathcal{C}^g \to \mathbb{R}^{\geq 0}$, and a function $\mathrm{FK} : \mathcal{C}^R \to \mathcal{C}^g$ that computes forward kinematics to the goal object's configuration space.

We compute the cost of a single primitive, a, applied to a state $x \in \mathcal{X}$ in two steps. First we compute the set $Q = \{q_1 \dots q_{p+1}\}$ of robot configurations achieved by the primitive. This can be obtained by forward propagating the controls in the primitive through the constraint f (Eq. 1). Then the cost of a primitive is:

$$\mathrm{c}_a(a, x) = \sum_{i=1}^{p} \mathrm{d}(\mathrm{FK}(q_i), \mathrm{FK}(q_{i+1})) \quad (4)$$

And the cost of a path, π:

$$\mathrm{c}_\pi(\pi, x_0) = \sum_{i=1}^{t} \mathrm{c}_a(a_i, x_i) \quad (5)$$

where x_i is the state reached by sequentially applying primitives $a_0 \dots a_{i-1}$ to x_0.

III. Heuristic Search-based Rearrangement Planning

Our goal is to perform an organized search across the high-dimensional state space. We make the following three assumptions on the planning instance:

Assumption 1: Contact between the robot and goal object, g, is restricted to the end-effector. We do allow contact between the full robot and all other objects in $\mathcal{M}$.

Assumption 2: Contact between objects is forbidden. For example, the robot cannot use one object to push another.

Assumption 3: All motions of, and interactions between, the robot and objects are quasistatic.

The search proceeds as follows. A list of vertices, V, each representing a state $x \in \mathcal{X}_{free}$ is maintained throughout the search. The list is initialized with the start state x_0. At each iteration of the search, a state is removed from the list and expanded. During expansion, a discrete set of primitives is applied. Each primitive is forward-propagated through a transition model that closely approximates the non-holonomic constraint f.

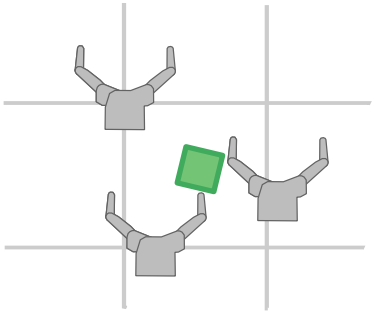

Fig. 3: A simple example of a hand pushing objects. The *basic* primitives allow the hand to translate along the grid lines. In this simple example, the object is "trapped", i.e. there are no primitives that allow it to move out of the cell.

The resulting states are added to V. The search ends when the state removed from V is a goal state.

The order that vertices are removed from V is determined by the particular search algorithm being used. Simple algorithms such as depth-first search or breadth-first search select states based on order of discovery. These search methods are unfocused and can lead to unnecessary expansion of several states unlikely to be on a path to the goal.

We would like to harness the power of heuristic-based graph search methods, such as A*. These select states based on the cost of the path to arrive at the state and the estimated remaining cost to achieve the goal. In addition, the algorithm maintains a list of expanded states and avoids unnecessary re-expansion when the same state is encountered along a different path.

In the following sections we define the four elements needed to use these heuristic search methods: (1) a transition model, (2) a discrete set of primitives, (3) a heuristic function and (4) a method for quickly determining whether a state has been expanded previously.

With these four elements defined we can use any heuristic search-based method to perform the search [2, 12, 19, 28]. Algorithm 1 outlines how our primitives, heuristics and mappings can be applied to an A*-based algorithm.

A. Transition Model

We require a transition model that closely approximates f (Eq. 1). We use a quasistatic pushing model with Coulomb friction [22]. In this model we assume pushing motions are slow enough that inertial forces are negligible; objects only move when pushed by the robot and objects stop immediately when forces cease to be imparted on the object.

B. Action Selection

We select a discrete set of actions, or primitives, to apply to each state that allow us to perform a feasible and focused search. We first select a set of primitives that

Algorithm 1 A heuristic search-based planner for solving rearrangement planning.

```
1: procedure SEARCH(x_0)
2:     B ← GetBasicPrimitives(U)
3:     v ← CreateVertex(x_0, 0, ComputeHeuristic(x_0))
4:     V.add(v)
5:     Q ← {}
6:     while not IsGoal(v.x) do
7:         for b ∈ B do
8:             ApplyPrimitive(b, v, Q, V)
9:         if not GoalObjectContact(v.x) then
10:            p ← GenerateContactPrimitive(v.x)
11:        else
12:            p ← GeneratePushPrimitive(v.x)
13:        ApplyPrimitive(p, v, Q, V)
14:        v ← Q.pop()
       return ExtractPath(v)
```

```
1: procedure APPLYPRIMITIVE(p, v, Q, V)
2:     x_new ← PhysicsPropagate(v.x, p)
3:     g_new ← v.g + c_a(p, v.x)
4:     if not V.find(x_new) then
5:         h ← ComputeHeuristic(x_new)
6:         v_new ← CreateVertex(x_new, g_new, h)
7:         Q.push(v_new)
8:         V.add(v_new)
9:     else
10:        v_new ← V.get(x_new)
11:        if g_new < v_new.g then
12:            Q.update(v_new, g_new)
       return v_new
```

move the robot without the explicit intent of creating contact or interaction with objects. These primitives are small motions of the robot defined by a coarse discretization of the control space. We label these primitives *basic* primitives.

These *basic* primitives are context agnostic; they are not specific to the rearrangement planning problem. We know contact with the goal object is critical to goal achievement. *Basic* primitives may achieve some contact with objects, but it is not guaranteed. Consider the simple example of a robot pushing an object in Fig.3. The *basic* primitives move in the four cardinal directions. The object is "trapped" , i.e. there is no primitive that can create enough contact to move it out of the current cell.

We could expand the set of *basic* primitives by discretizing the control space more finely. However the size of the primitive set is directly related to the branching factor of the search, so any large increase affects the computation time.

Instead we propose to augment the primitive set applied at each state with a dynamically-generated action aimed at creating or maintaining contact with the goal object. We label these primitives as *contact* and *pushing* primitives, respectively. The use of these primitives focuses our search to areas of the full state space that are likely to lead to the goal.

We observe that it is often possible to solve the 2PTBVP in the lower-dimensional subspace containing only the robot. We use this to design the *contact* and *pushing* primitives.

To generate a *contact* primitive to be applied to a state $x \in \mathcal{X}_{free}$, we first find a pose, q_{con}, of the robot such that the end-effector is in or near contact with the goal object in configuration $o^g \in x$. Then we solve the 2PTBVP in the robot configuration space to generate motions that move from q to q_{con}. During the search, we apply a *contact* primitive to any state where the robot and goal object are not in or near contact.

A *pushing* primitive aims to push the goal object toward the goal region. During the search we apply a *pushing* primitive to any state where the robot and goal are in or near contact. To generate the primitive we produce any motion of the robot that moves in the direction of the goal object. We detail a specific example *contact* and *pushing* primitive in our experiments in Section IV.

The *basic* and *contact* primitives are similar to transit actions defined by Simeon et al. [31], where the robot changes configuration without moving an object. The *pushing* primitive mimics transfer actions, where an object is grasped and reconfigured. We note that the correspondence is not exact, *basic* and *contact* primitives may inadvertently make contact with objects in the scene. Likewise, *pushing* primitives may lose contact with the goal object in the middle of primitive execution.

C. Heuristic

Developing a heuristic that estimates the distance between a state and the goal region is challenging because the goal is underspecified. In particular, the goal is defined only by the configuration of g. The configuration of the robot is not defined. As a result, common robot-configuration based heuristics are not directly applicable.

However, two observations of the problem can be used to generate a useful heuristic. First, by definition of the problem, contact with the goal object is required for goal achievement. Due to Assumptions 1 and 2, this contact must be between the end-effector and the object. Second, the robot must stay in contact with the goal object for the object to move, due to Assumption 3.

Using these observations we develop a two part heuristic to estimate the distance between a state $x \in$

$\mathcal{X}_{free}$ and $\mathcal{X}_G$:

$$h(x) = \hat{\mathrm{d}}_{con}(x) + \tag{6}$$
$$\hat{\mathrm{d}}_{move}(x) \tag{7}$$

Eq. 6 estimates the distance to make contact with the goal object. Eq. 7 estimates the distance the end-effector must move to push the goal object to the goal region.

Distance to contact: We compute $\hat{\mathrm{d}}_{con}$ by approximating the end-effector with the smallest enclosing sphere. If this sphere penetrates the object, $\hat{\mathrm{d}}_{con} = 0$. Otherwise, $\hat{\mathrm{d}}_{con}$ is the translational distance between the closest points on the sphere and object under the metric d.

Proposition $\hat{\mathrm{d}}_{con}$ is a lower bound on the true cost to make contact with the goal object.

Proof: Approximating the end-effector pose with the sphere means all rotations of the end-effector have $\hat{\mathrm{d}}_{con} = 0$. Thus our estimate of the rotation distance is a lower bound of the true distance. The shortest translational distance the end-effector can move to make contact is the distance between the two closest points on the end-effector and object. Selecting the closest point on the sphere to the object ensures we underestimate this distance. Since we underestimate translational and rotational distance, we must underestimate the true distance. ■

Distance to goal: We compute $\hat{\mathrm{d}}_{move} = \mathrm{d}(o^g, p_{goal}) - r_{goal}$. This is the straight line distance from the object location to the edge of the goal region.

Proposition $\hat{\mathrm{d}}_{move}$ is a lower bound on the true cost to move the object to the goal.

Proof: $\hat{\mathrm{d}}_{move}$ is the shortest distance the object can move and still achieve the goal. By the quasistatic assumption, contact must be maintained between robot and object for the object to move. As a result, $\hat{\mathrm{d}}_{move}$ must also be the shortest distance the robot could move. ■

D. State equality

To more effectively use heuristic-based search methods we need to recognize when the search has encountered the same state along two different paths. One commonly used method is to define a mapping from a state $x \in \mathcal{X}_{free}$ to a discrete bin. Any two states that map to the same discrete bin are considered equivalent.

This method relies on the assumption that two states that start near each other will end near each other when a primitive is applied. The discontinuous nature of pushing interactions means this assumption often does not hold. Consider the example from Fig.4. A very small change in initial object pose (Fig.4a) leads to large difference of final pose (Fig.4b), despite the same primitive being applied. As a result we can only mark two states as equivalent when they are exactly equivalent. To do this, we track all states encountered along the search and use a Geometric Near-neighbor Access Tree (GNAT) [5] to efficiently find the nearest neighbor to a new state and check equality.

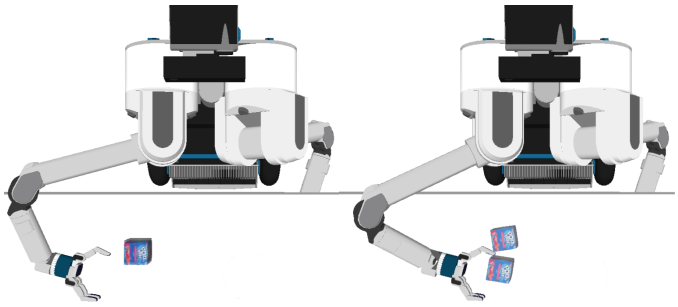

(a) Two states with object pose that differs by less than 1cm

(b) After executing the same action, the final object pose differs significantly.

Fig. 4: Under typical mappings from continuous to discrete states, the two states depicted in (a) would be labeled the same. However, applying the same primitive leads to very different final states (b).

IV. Experiments and Results

We have implemented our planner using the Boost Graph Library [1]. Using the planner we test three hypotheses:

H1: The use of *contact* and *pushing* primitives improves success rate when compared to a planner that uses only *basic* primitives.

H2: The use of the search-based planner is able to produce more optimal paths than current randomized methods.

H3: The goal-directed nature of the *contact* and *pushing* primitives allows the search planner to find solutions in plan times comparable to current randomized methods.

The experiments are conducted using a 7 degree-of-freedom arm, first in simulation and then on the real robot.

A. Experiment Setup

We evaluate our hypotheses on 12 scenes, each containing between 1 and 7 movable objects. Fig.6 shows an example scene. We use the same robot start configuration and goal region in all 12 scenes. Each scene contains the same goal object, but its start location, and the number and start location of all other objects differs across scenes. The objects are placed on a table. The table is treated as an obstacle that the robot is forbidden to contact. We set $r_{goal} = 10$cm for all scenes. We use a quasistatic model of physics to propagate interactions between the objects and the robot. We model only planar

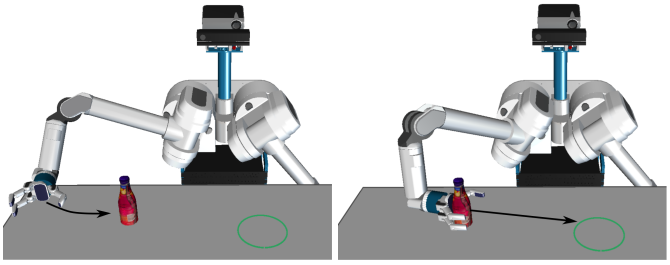

(a) The *contact* primitive tries to make contact between the palm and the goal object.

(b) The *pushing* primitive moves the end-effector in the direction of the goal region

Fig. 5: Context specific primitives are used to guide the search.

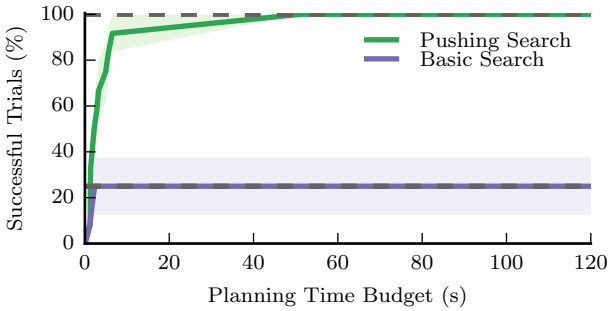

Fig. 7: Percent of successfully solved scenes as a function of plan time. Shaded regions represent standard error.

pushing actions (no toppling) for the 7-DOF arm. Thus $\mathcal{C}^R = \mathbb{R}^7$ and $\mathcal{C}^1 = \cdots = \mathcal{C}^m = SE(2)$. We follow the ideas from [16] and constrain the motion of the end-effector to the plane parallel to the pushing support surface. As required by Assumption 1 we invalidate any robot motions that create contact between the goal object and any part of the robot other than the manipulator. We note that we allow and model pushing contact between the full arm and all objects except the goal object. As required by Assumption 2, we also invalidate any robot motions that lead objects to contact one another.

The planner is given a set of 6 *basic* primitives. Each of these six actions moves the end-effector in $SE(2)$ at a pre-defined maximum velocity (positive or negative) along a single axis (x, y, θ). We use a Jacobian psuedo-inverse controller to compute the 7-DOF velocities that achieve the desired end-effector motion. The maximum linear (0.5 m/s) and angular (1.0 rad/s) velocity limits were selected to ensure the resulting 7-DOF joint velocities remained within safety limits defined for the robot. We set the duration of each *basic* primitive to 0.2 s.

The *contact* primitive defined on an object moves the end-effector along a straight line in $SE(2)$ toward the goal object until either contact is made or an invalid state is encountered. During the motion, the end-effector is rotated to face the palm of the hand to the object centroid. The *pushing* primitive moves the end-effector in a straight line in $SE(2)$ along the vector normal to the palm. Fig.5 depicts an example *contact* and *pushing* primitive.

1) Baseline Planners: We compare the performance of our search-based planner (denoted ***Pushing Search*** in all results) against three baseline planners. First we compare against a planner that uses only the *basic* primitives defined above. We denote this planner as ***Basic Search*** in all results and discussions.

Next, we compare against the RRT planner used in [16]. This planner grows a search tree from the start state toward a goal region by iteratively sampling a random state from $\mathcal{X}_{free}$, and growing the tree toward the sample. During tree extension we sample $k = 3$ actions and use the quasistatic model of physics to forward-simulate each action. We keep the action that ends closest to the sample under a weighted euclidean distance metric. We denote this planner as ***Pushing RRT*** in all results and discussions.

Finally, we compare against an altered version of the RRT that includes manipulation primitives similar to our *contact* and *pushing* primitives. In this planner, during tree extension, a *contact* primitive is applied to move the robot near an object, then a *pushing* primitive is applied to move the object. The moved object and push direction are indicated in the sampled state the tree is growing towards. We denote this planner as ***Primitive RRT*** in all results and discussions.

For the RRT-based planners, we do not apply Assumption 1 and allow contact between the full robot and the goal object. Our quasistatic physics model does not model object-to-object interaction. As a result we must still impose Assumption 2 on the RRT planners.

B. Analysis

1) Effect of Manipulation Primitives: We first compare the performance of the ***Pushing Search*** planner against the ***Basic Search*** planner. We run each planner on the 12 scenes. We allow the planners 300s to find a solution for each scene. If a solution is not returned within this time the run is marked failure. We use a Weighted A* [28] search with $w = 5.0$ in all searches.

Fig.7 shows the percent of scenes solved successfully as a function of plan time. When using only *basic* primitives, the search-based planner is only able to solve 3 of the 12 scenes within the time limit. The addition of the *contact* and *pushing* primitives allows the planner to solve all 12 scenes in the 300s time limit. This supports **H1**: **The use of *contact* and *pushing* actions improves success rate of search-based planners.**

2) Path Optimality: Next we compare the optimality of the paths produced by the ***Pushing Search*** planner against the paths produced by the ***Pushing RRT*** and the ***Primitive RRT*** . We run the randomized planners 30

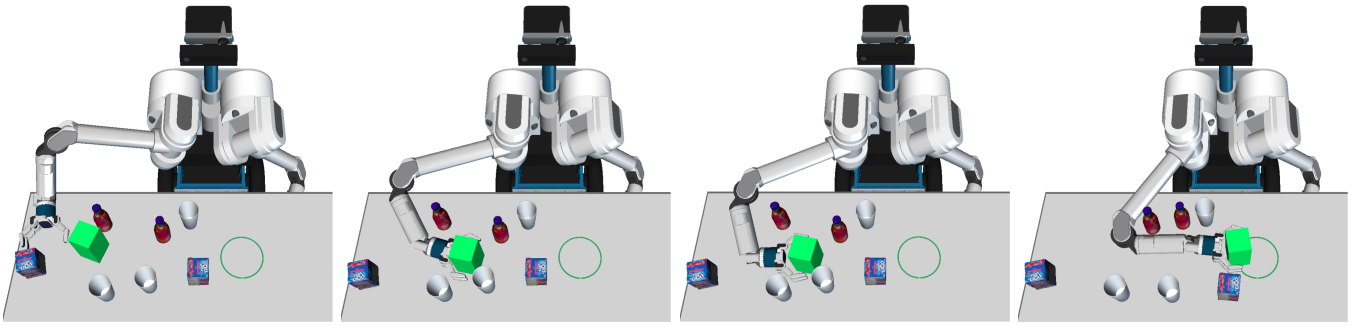

(a) In many paths generated by the ***Pushing Search*** , the robot's initial motion is a *contact* primitive.

(b) *Basic* primitives are used to reposition the robot when stuck. Here the box cannot be pushed further without contacting the glass.

(c) After the *basic* rotation primitive reorients the arm, the robot can continue a *pushing* primitive towards the goal.

(d) As the object is pushed to the goal, the fingers and upper arm are used to push clutter aside.

Fig. 6: A solution for one of the 12 simulation scenes. The robot is tasked with pushing the green box into the green circle.

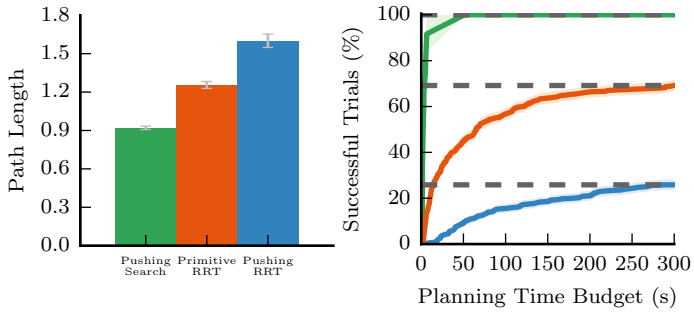

(a) A comparison of average path length across all scenes.

(b) The percent of successful trials as a function of plan time. Shaded regions represent standard error.

Fig. 8: Performance of the ***Pushing RRT*** as compared to RRT based planners that solve rearrangement problems

times on each scene, allowing each planner up to 300s to find a solution. We compute the path length as the average across all successfully planned paths.

Fig.8a shows a comparison of the average path length for the three planners. A one-way ANOVA with Tukey HSD post-hoc analysis reveals that the path length for paths produced by our ***Pushing Search*** algorithm is significantly smaller that the average path lengths for paths produced by the ***Pushing RRT*** ($p = 0.00$) and the ***Primitive RRT*** ($p = 0.00$). This supports our second hypothesis: **H2**: **Our search-based planner with *contact* and *pushing* primitives is able to produce more optimal paths than randomized planners.**

3) Plan Time: Fig.8b compares the percent of successful trials as a function of plan time for the ***Pushing Search*** planner to that of the ***Pushing RRT*** and ***Primitive RRT*** . As can be seen, the ***Pushing Search*** is able to solve more trials faster than either RRT method. This supports our third hypothesis: **H3**: **The use of goal-focused primitives allows the search based planner to find solutions faster than the randomized planners.**

These results may appear surprising because randomized planners, and in particular bi-directional randomized planners, have been shown to be faster than search-based methods in high-dimensional spaces. Because we cannot solve the 2PTBVP, both the ***Pushing RRT*** and ***Primitive RRT*** must be implemented as single-directional planners where only one tree is grown rooted at the start state. Additionally, extensions to the trees fail to extend all the way to the sampled state, especially in the ***Pushing RRT*** . As a result, the planners lack the ability to perform focused extension toward the goal. Conversely, the *contact* and *pushing* are explicitly designed to lead the search toward the goal. We believe this accounts for the fast planning times exhibited by the planner on our test scenes.

C. Real Robot Experiments

Next we perform a set of trials on the real robot. We use the planners to generate plans for 12 scenes similar to those used in simulation. For each scene we randomly place between 1 and 6 objects on a table within the reachable workspace of the robot. The pose of these objects is detected using AprilTag fiducial markers [26]. Fig.1, Fig.2 and Fig.9 depict 3 of the 12 scenes.

We use the same parameters for the ***Pushing Search*** algorithm as used in the simulation results. All 12 scenes have the same start configuration for the robot and goal region. We set $r_{goal} = 10$cm across all scenes.

For each trial we record the time to generate a plan using our planner. Fig.10b shows the percent of scenes solved as a function of plan time. The planner is able to successfully find solutions for 11 of the 12 scenes.

D. Qualitative Analysis

Comparing the paths across the 24 scenes (12 simulation, 12 real robot) reveals some interesting qualitative

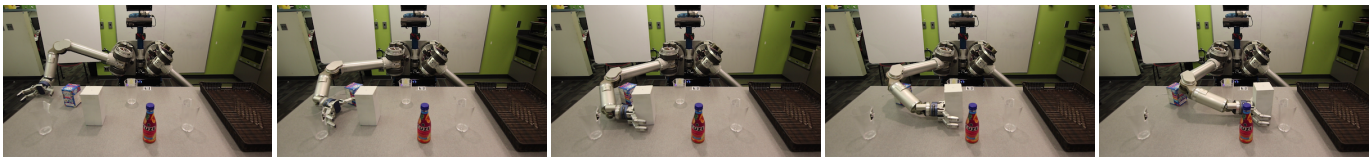

Fig. 9: The ***Pushing Search*** planner is able to find a path using a single *contact* primitive followed by a single *pushing* primitive.

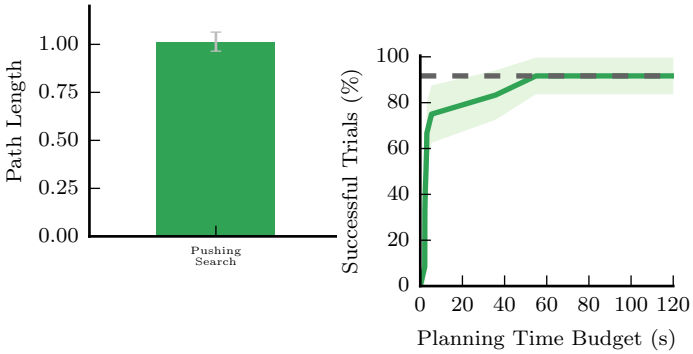

(a) The lengths of the paths generated by the ***Pushing Search*** planner.

(b) The percent scenes successfully solved by the ***Pushing Search*** planner as a function of plan time. Shaded regions represent standard error.

Fig. 10: Results from trials on 12 scenes for the real robot.

aspects of the solutions. First, the *contact* and *pushing* primitives comprise the majority of the paths. This is not surprising as they have the most effect on the heuristic cost. In particular, a successfully applied *contact* primitive reduces Eq. 6 to zero while the *pushing* primitive reduces Eq. 7. The result is that the planner is guided to create paths that move directly to the goal object then push it almost directly to the goal (Fig.6, Fig.9), using the fingertips, back of the hand and arm to clear clutter along the way.

In these solutions, we see *basic* primitives used as "regrasp" motions. Consider the motion between Fig.6b and Fig.6c. In Fig.6b the green box cannot be pushed further without hitting the glass, an action that violates Assumption 2 . A *basic* rotation primitive is used to reorient the robot, allowing further pushing of the object.

V. Discussion

In this work, we formulate a method for applying heuristic search to rearrangement planning by pushing. We show the use of dynamically generated, problem specific primitives that are critical to goal achievement. These primitives, combined with an informative and admissible heuristic, guide the search to promising areas of state space. Our experiments show we are able to quickly produce low cost paths for several problems.

While our experiments are promising, our formulation imposes several limitations to address in future work. We define the cost function we optimize in the configuration space of the goal object. This simplifies the definition and computation of $\hat{d}_{con}$ (Eq. 6) and $\hat{d}_{move}$ (Eq. 7). However, it is often more desirable to express cost in the configuration space of the robot, e.g. path length in joint space. For manipulators such as the one used for our experiments, it is difficult to compute meaningful and admissible estimates of the distance in joint space to contact the goal object and the minimum joint motions that move the object. However, if we could find a computationally tractable method for estimating these values, we could not only use a more desirable cost function, but we could also eliminate Assumption 1 . This would allow solutions that exhibit whole-arm or even whole-body interactions with the goal object.

Assumption 2 prevents solutions where the robot uses one object to push another. This is particularly restrictive when the robot is working in highly cluttered spaces. We could eliminate this assumption at the expensive of a less informed heuristic. In particular, we could define $\hat{d}_{con}$ as the distance to the *nearest* object rather than distance to the goal object. This implies that the robot can push the nearest object and create a chain of pushes that moves the goal object. This will often grossly underestimate the true distance to goal. Future work should examine whether a tighter bound could be inferred from the structure of the state.

While the quasistatic assumption is common in planning pushing interactions [3, 8, 20–24, 27], incorporating dynamic interactions increases the space of possible solutions. For example, the robot could simply move to the object, strike it and remain in place as the object slides to the goal [14, 36]. A naive incorporation of dynamic interactions into the heuristic sets $\hat{d}_{move} = 0$. This is a looser bound than we are able to formulate under the quasistatic assumption. Additionally, we incur the penalty of a state space that doubles in size, as we must track the configurations *and* velocities of all objects.

Despite these limitations we believe the results presented here are promising. The methods for selecting a discrete set of primitives from the continuous space of robot motions can be applied to online planning and robust path generation for the rearrangement problem. This will allow planners to represent and reason about uncertainty and incorporate feedback, resulting in better execution of rearrangement plans.

Acknowledgments

References

[1] *The Boost Graph Library: User Guide and Reference Manual*. Addison-Wesley Longman Publishing Co., Inc., 2002. ISBN 0-201-72914-8.

[2] S. Aine, S. Swaminathan, V. Narayanan, V. Hwang, and M. Likhachev. Multi-heuristic A*. In *RSS*, 2014.

[3] S. Akella and M. T. Mason. Posing polygonal objects in the plane by pushing. In *IEEE ICRA*, 1992.

[4] J. Barry, K. Hsiao, L. P. Kaelbling, and T. Lozano-Pérez. Manipulation with multiple action types. In *ISER*, 2012.

[5] S. Brin. Near neighbor search in large metric spaces. In *Proceedings of the 21st International Conference on Very Large Data Bases*, 1995.

[6] P.C. Chen and Y.K Hwang. SANDROS: A dynamic graph search algorithm for motion planning. In *IEEE ICRA*, 1992.

[7] B. Cohen, S. Chitta, and M. Likhachev. Search-based planning for manipulation with motion primitives. In *IEEE ICRA*, 2010.

[8] M.R. Dogar and S.S. Srinivasa. Push-grasping with dexterous hands: Mechanics and a method. In *IEEE/RSJ IROS*, 2010.

[9] M.R. Dogar and S.S. Srinivasa. A framework for push-grasping in clutter. In *RSS*, 2011.

[10] J.D. Gammell, S. Srinivasa, and T.D. Barfoot. Informed RRT*: Optimal sampling-based path planning focused via direct sampling of an admissible ellipsoidal heuristic. In *IEEE/RSJ IROS*, 2014.

[11] J.D. Gammell, S. Srinivasa, and T.D. Barfoot. Batch informed trees (BIT*): Sampling-based optimal planning via the heuristically guided search of implicit random geometric graphs. In *IEEE ICRA*, 2015.

[12] P.E. Hart, N.J. Nillson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. In *TSSC*, 1968.

[13] K. Hauser and V. Ng-Thow-Hing. Fast smoothing of manipulator trajectories using optimal bounded-acceleration shortcuts. In *IEEE ICRA*, 2010.

[14] J.A. Haustein, J.E. King, S.S. Srinivasa, and T. Asfour. Kinodynamic randomized rearrangement planning via dynamic transitions between statically stable configurations. In *IEEE ICRA*, 2015.

[15] S. Karaman and E. Frazzoli. Sampling-based algorithms for optimal motion planning. *IJRR*, 30(7): 846–894, 2011.

[16] J.E. King, J.A. Haustein, S.S. Srinivasa, and T. Asfour. Nonprehensile whole arm rearrangement planning on physics manifolds. In *IEEE ICRA*, 2015.

[17] S.M. LaValle. Rapidly-exploring Random Trees: A new tool for path planning. 1998.

[18] M. Likhachev and D. Ferguson. Planning long dynamically-feasible maneuvers for autonomous vehicles. In *RSS*, 2008.

[19] M. Likhachev, G. Gordon, and S. Thrun. ARA*: Anytime A* with provable bounds on sub-optimality. In *NIPS*, 2004.

[20] K. Lynch. Locally controllable polygons by stable pushing. In *IEEE ICRA*, 1997.

[21] K. Lynch and M.T. Mason. Controllability of pushing. In *IEEE ICRA*, 1995.

[22] K. Lynch and M.T. Mason. Stable pushing: Mechanics, controllability, and planning. In *WAFR*, 1995.

[23] M.T. Mason. Mechanics and planning of manipulator pushing operations. *IJRR*, 5(3):53–71, 1986.

[24] D. Nieuwenhuisen, A. Frank van der Stappen, and M.H. Overmars. Path planning for pushing a disk using compliance. In *IEEE/RSJ IROS*, 2005.

[25] D. Nieuwenhuisen, A. Stappen., and M. Overmars. An effective framework for path planning amidst movable obstacles. In *WAFR*, 2008.

[26] E. Olson. AprilTag: A robust and flexible visual fiducial system. In *IEEE ICRA*, 2011.

[27] M.A. Peshkin and A.C. Sanderson. Planning robotic manipulation strategies for workpieces that slide. *IJRA*, 4(5):524–531, 1988.

[28] I. Pohl. Heuristic search viewed as path finding in a graph. *Artificial Intelligence*, 1(3):193–204, 1970.

[29] G. Sánchez and J-C. Latombe. On delaying collision checking in PRM planning - application to multi-robot coordination. *IJRR*, 21(1):5–26, 2002.

[30] S. Sekhavat, P. Švestka, J.-P. Laumond, and M.H. Overmars. Multi-level path planning for nonholonomic robots using semi-holonomic subsystems. *IJRR*, 17(8):840–857, 1998.

[31] T. Siméon, J-P. Laumond, J. Cortés, and A. Sahbani. Manipulation planning with probabilistic roadmaps. *IJRR*, 23(7–8):729–746, 2004.

[32] M. Stilman. and J. Kuffner. Navigation among movable obstacles: Real-time reasoning in complex environments. In *IEEE-RAS Humanoids*, 2004.

[33] M. Stilman, J. Schamburek, J. Kuffner, and T. Asfour. Manipulation planning among movable obstacles. In *IEEE ICRA*, 2007.

[34] J. van den Berg, M. Stilman, J. Kuffner, M. Lin, and D. Manocha. Path planning among movable obstacles: a probabilistically complete approach. In *WAFR*, 2008.

[35] D. J. Webb and J. van den Berg. Kinodynamic rrt*: Optimal motion planning for systems with linear differential constraints. *CoRR*, 2012.

[36] S. Zickler and M. Velosa. Efficient physics-based planning: sampling search via non-deterministic tactics and skills. In *AAMAS*, 2009.